Fault-tolerant quantum communication with rare-earth elements and superconducting circuits

Ashley M. Stephens,[1,*] Jingjing Huang,[2,3] Kae Nemoto,[1] and William J. Munro[2,1]

[1]*National Institute for Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan*
[2]*NTT Basic Research Laboratories, 3-1 Morinosato-Wakamiya, Atsugi, Kanagawa 243-0198, Japan*
[3]*California Institute of Technology, Pasadena, California 91225, USA*

(Dated: September 19, 2012)

We present a scheme for quantum communication that involves the distribution of a topological cluster state throughout a quantum network. Photon loss and other errors are suppressed by optical multiplexing and entanglement purification. We discuss implementation of an individual node composed of Erbium spins (single atom or ensemble) coupled via flux qubits to a superconducting resonator, allowing for deterministic local gates, stable quantum memories, and emission of photons in the telecom regime.

PACS numbers: 03.67.Hk, 03.67.Pp

Quantum communication is predicated on the ability to quickly and reliably entangle two or more quantum systems that are separated by geographically large distances [1]. What makes quantum communication difficult is that photons are easily lost during transmission, due to attenuation in optical fibers and inefficient optical coupling. Although photon loss is expected to predominate, other sources of error such as decoherence and imprecise quantum control may ultimately limit the entanglement fidelity. Schemes for quantum communication should account for all such sources of error without placing unreasonable demands on the complexity or accuracy of the hardware. This property is also desirable for distributed quantum computation, in which a number of small arrays of qubits are connected to form a quantum computer that is large enough to solve interesting problems.

It has long been known that reliable quantum communication is possible using a network of relatively simple devices known as quantum repeaters, analogous to how optical amplifiers are used for classical communication [2–4]. The communication channel is divided into short segments, pairs of qubits in adjacent nodes are entangled with high fidelity [5–7], and the range of entanglement is extended to the endmost nodes [8]. The basic elements of such a scheme have been demonstrated in the laboratory [9–12]. However, in a practical setting, the coherence time of the quantum memories in the system will limit the communication distance and the accuracy of the quantum gates at the nodes will limit the entanglement fidelity [13]. Recently, a number of alternatives to the orthodox repeater network have been considered, including schemes based on graph states [14, 15] and schemes that incorporate fault-tolerant error correction [16–18]. Error correction could, for example, extend the effective lifetime of quantum memories during communication.

Quantum communication will soon scale beyond the laboratory, but theoretical and experimental work on fundamental components should be informed by coherent and comprehensive designs for schemes for quantum communication that account for all sources of error. In particular, if a scheme for quantum communication adheres to the principles of fault tolerance, efficient and reliable communication over arbitrarily long distances is possible, but only if photon loss and other errors are below certain thresholds, as per the threshold theorem of quantum computing [19]. As such schemes emerge, it is important to identify physical systems that meet the requirements of scalability, to establish threshold error rates for the fundamental components, and to understand the associated performance and resource requirements.

Here, we study a scheme for quantum communication that combines aspects of traditional repeater networks with fault-tolerant error correction. The foundation of the scheme is optical multiplexing, which reduces the effective loss rate between nodes [20]. Then, purification is used to increase the fidelity of entanglement between nodes [13]. Combining these techniques significantly lessens the requirements on the communication channel. Finally, topological cluster-state error correction [21, 22] is applied to ensure that communication over large distances is reliable. Topological cluster-state error correction requires low connectivity but exhibits a relatively high threshold [23], making it suitable for implementation in a variety of physical systems. After accounting for photon loss and errors arising from the communication channel ($\sim$80% and $\sim$10% per qubit, respectively) and errors arising from local gates ($\sim$0.1% per gate), we find that the scheme is scalable to large distances at an end-to-end rate of $\sim$kHz with about fifty qubits per node. We discuss implementation of an individual node composed of Erbium spins (single atom or ensemble) coupled via flux qubits to a superconducting resonator.

Physical details— In general, a scheme for quantum communication requires a way to transmit quantum information between nodes and a way to process that same information at the nodes. Hybrid systems, which marry photonic qubits with matter qubits, are natural candidates to consider [24]. Here, we consider an emitter-receiver model, in which a matter qubit with an optical transition is placed in a cavity coupled to an optical fiber. Our scheme requires quantum memories that are stable over multiples of the two-way time of flight between adjacent nodes. Also, it requires deterministic local gates and transmission in the telecom regime.

To meet these requirements is challenging, so we turn our attention to an unusual candidate. Superconducting systems have emerged as promising devices for quantum computing [25]. Here, we propose that each node contain a number of

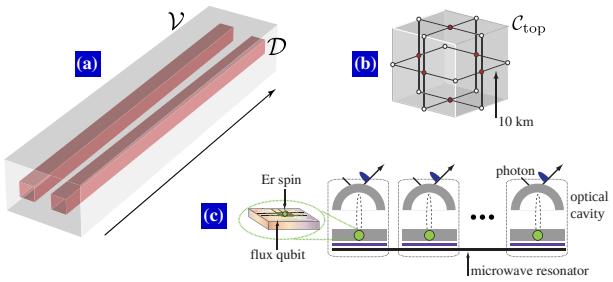

FIG. 1. Scheme for quantum communication. (a) Communication is achieved by transmitting logical qubits through a topological cluster state. Logical qubits are defined by the region $\mathcal{D}$, wherein qubits are measured in the Z basis. Qubits in the region $\mathcal{V}$ are measured in the X basis to enable error correction. (b) The connectivity of the network mimics the structure of the topological cluster state, $\mathcal{C}_{\rm top}$. There is one cluster-state qubit per node, and gates between qubits in adjacent nodes are performed using purified entangled pairs. (c) The architecture of the nodes features Erbium spins, where local gates are performed using superconducting qubits and resonators.

gap-tuneable flux qubits [26, 27] coupled to a resonator (LC resonator, high-Q superconducting coplanar resonator) [28]. Each qubit can be brought into and out of resonance with the resonator as required. The Jaynes-Cummings interaction permits an iSWAP interaction between the flux qubit and the resonator. A sequence of operations permits a deterministic entangling gate between two flux qubits in the same node. Josephson bifurcation amplifiers can be used to perform fast, high-fidelity, single-shot measurements [29]. One of the advantages of this approach is that the flux qubits can be far detuned from the resonator, effectively turning off the interaction and avoiding significant crosstalk. We are, however, limited to having only one flux qubit interacting with the resonator at a time. Sequential interactions might limit the scalability of the system for some applications in quantum computing, but our requirements are less stringent. The duration of an entangling gate is expected to be 150 ns.

Aside from the ability to perform local gates, we need a way to couple qubits in adjacent nodes—that is, we need to integrate the flux qubits with the optical cavity qubits. Recently, strong coherent coupling between a flux qubit and an ensemble of electron spins has been demonstrated with nitrogen-vacancy (NV) centers in diamond [30]. The electron spin ensemble couples to the flux qubit at 2.88 GHz. NV centers in diamond also have an optical transition, but at 637 nm it is outside the telecom band. However, coupling between a rare earth element and a microwave resonator has also been demonstrated [31, 32]. Er^{3+} spin qubits (either one or an ensemble) doped in a Y_2SiO_5 crystal can couple to the flux qubit-resonator system using a magnetic field to induce energy splitting. The flux qubit-resonator system can then be used to perform local gates between different Er^{3+} qubits, which will serve as quantum memories in our scheme. Furthermore, Er^{3+} has a ${}^4I_{I5/2} - {}^4I_{I3/2}$ transition in the telecom C-band (at $\sim$1540 nm) allowing coupling between an optical field and the electron spin as required [33].

Quantum multiplexing— With a hybrid system at each node, we can generate entanglement between a pair of cavity qubits in neighboring nodes in a simple way [20]. First, we prepare both qubits in the state $(|0\rangle + |1\rangle)/\sqrt{2}$. Then, we couple a single photon to both qubits and send it to a detector. Detection heralds an entangled state $(|01\rangle + |10\rangle)/\sqrt{2}$. In practice, if the node separation is 10 km, loss due to attenuation and inefficient sources, detectors, and coupling is expected to exceed 80%. Simply repeating until success is, in principle, sufficient for quantum communication and distributed quantum computation [2, 3]. However, every additional attempt extends the time over which quantum memories are required to be stable and slows the rate of communication.

Instead, multiplexing can reduce the effective loss rate between nodes with a modest increase in complexity [20, 34–37]. At the first node we couple a number of photons with a number of qubits, effectively initiating several attempts to generate entanglement at once. Then, these photons are temporally multiplexed and transmitted to the second node, preceded by a classical herald. When the herald is received, the first photon is coupled with a qubit at that node. As before, the photon is sent to a detector. However, in this case, if the photon is lost, the qubit is re-prepared in time for the second photon, and so on until a success is reported. Once that qubit is entangled the remaining photons are sent to another qubit, and so on until all incoming photons have been depleted [52].

It is important that multiplexing be tolerant of the various errors that can arise. By design, photon loss is tolerated. An error during preparation of the cavity qubits or an error due to decoherence of the quantum memories during transmission of the photons will simply lower the fidelity of the entanglement in the event that the attempt succeeds. Such an error cannot affect more than one pair of qubits that is eventually accepted, so, for these errors, attempts are effectively independent. Measurement errors might not be so benign. If an attempt to generate entanglement is successful but the photon is not detected due to a measurement error, then the pair of qubits is rejected as if the photon was lost. However, in the case of a dark count, we might accept a pair of qubits that are not, in fact, entangled. Thus, we require single-photon detectors that are reliable enough so that the cumulative probability of a dark count during a multiplexed series of attempts does not limit the fidelity of the entanglement. Recent results suggest that this can be achieved [38]. Ultimately, we are left with a number of entangled pairs of fidelity F shared between adjacent nodes in the network.

Entanglement purification— Next, to increase the fidelity of entanglement between adjacent nodes, we turn to purification [5–7]. By executing a bipartite circuit, a number of entangled pairs can be purified to a single pair of higher fidelity, conditional upon a set of measurement results. Purification requires classical communication between the two nodes to coordinate the conditioning. Successive rounds of purification can be performed, using already purified pairs, until sufficient fidelity is achieved. This process can be structured in various ways in order to, for example, maximize the rate at which entan-

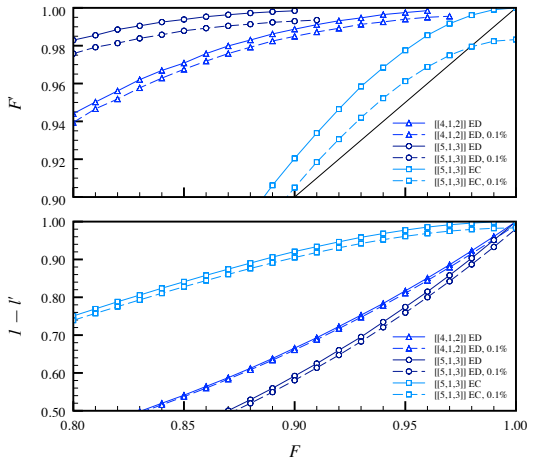

FIG. 2. Output fidelity, F', and reliability, $1-l'$, as a function of input fidelity, F, for purification based in error-detection (ED) and error-correction modes (EC), where l' is the probability that purification fails to produce an entangled pair. Solid lines are for perfect local gates and dashed lines are for local gates with an error rate of 0.1%.

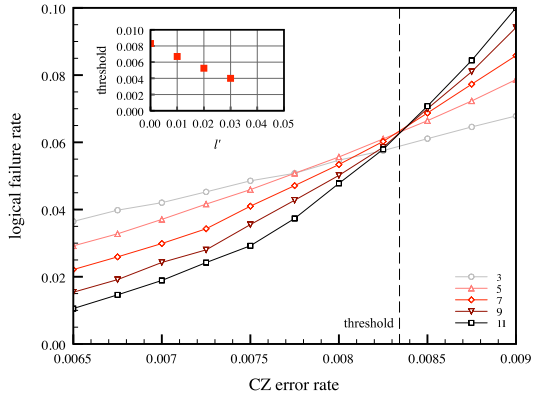

FIG. 3. Results of numerical simulation of topological error correction for various code distances, d. A lattice of $d \times d \times d$ unit cells is associated with a distance-d code that may fail for $\geq \lceil d/2 \rceil$ errors. The inset shows the threshold as a function of the fraction of discarded CZ gates during preparation, l'.

glement is purified or minimize the number of pairs that are consumed [39]. Here, we will consider a variation of purification based on Calderbank-Shor-Steane error-correction codes [13, 40]. For an $[[n, k, d]]$ code, n entangled pairs are required for each round. For codes with $d > 2$, we can attempt to correct errors (error-correction mode) instead of simply detecting them (error-detection mode). Figure 2 shows the performance of purification for various codes in both modes. At this point we make the assumption that the error rate of the local gates is 0.1%. For an initial fidelity $F = 0.900$, two rounds of the [[4,1,2]] code results in a final fidelity $F' = 0.997$, while one round of the [[5,1,3]] results in a final fidelity $F' = 0.993$. This is a greater increase in fidelity per round than for standard two-qubit purification.

Topological error correction— With high-fidelity entanglement between adjacent nodes, we want to efficiently establish entanglement of arbitrary fidelity that spans the entire network. To do this, we use fault-tolerant error correction. Of the many schemes for error correction, topological cluster-state error correction is particularly promising, since it tolerates a relatively high rate of physical errors while being local in a three-dimensional setting. [21, 22]. The main ingredient of the scheme is the topological cluster state shown in Fig. 1(b). One way to prepare such a state is to initialize each qubit in the $|+\rangle$ state then to apply CZ gates between neighboring qubits. In our case, the topological cluster state is prepared with one qubit per node. CZ gates are executed using the entangled pairs shared between adjacent nodes [41].

Once the state is prepared, communication proceeds with a sequence of single-qubit measurements. The state is divided into regions that determine the appropriate measurement basis. Logical qubits are defined by regions of the state measured in the Z basis. It is these logical qubits that are transmitted from one end of the network to the other. The rest of the qubits are measured in the X basis to obtain a syndrome. Error correction involves finding the most likely set of errors that is consistent with this syndrome [21–23]. For error rates below a certain threshold value, increasing the distance of the code by increasing the extent of the topological cluster state will decrease the likelihood of logical errors.

Since the cluster state is distributed throughout the network, the CZ gates rely on the availability of entangled pairs. It is possible to allocate enough resources when generating entangled pairs to ensure that the probability that a CZ cannot be performed is lower than its error rate. Then, errors due to unreliable CZ gates are managed like unknown errors. However, because they are heralded, these unreliable gates are easier to manage [23]. If an entangled pair is unavailable, the CZ gate is discarded and the associated qubits are treated as if they were lost. Then, the lattice is deformed to avoid the lost qubits and error correction proceeds. This introduces a trade-off between the fraction of CZ gates that are discarded and the threshold error rate of the remaining gates [23].

We simulate a topological cluster state with periodic boundary conditions in all three dimensions [53]. Figure 3 shows the logical failure rate as a function of the CZ error rate for various code distances, where we assume that the error rate of preparation and measurement of qubits at the nodes is 0.1% and that CZ gates between nodes are completely reliable. We observe a threshold error rate of approximately 0.83%. The inset to Fig. 3 shows the expected tradeoff between the fraction of discarded gates and the threshold. In Ref. [23], the threshold for lost qubits is found to be 24.9%, equal to the threshold for bond percolation on a cubic lattice [42]. Here, the corresponding threshold for the fraction of discarded gates is approximately 5%. The threshold in Ref. [23] is higher because loss is assumed to occur only at the beginning (or end) of preparation. In our case, qubits are effectively lost (due to discarded gates) at intermediate times.

Performance and overhead— To determine the overhead of

F	0.784	0.816	0.835	0.872	0.907	0.919	0.930	0.951	0.963
p_{local}	0.0005	0.001	0.0005	0.0005	0.0008	0.0005	0.0005	0.001	0.0005
l'	0.008	0.008	0.02	0.008	0.008	0.02	0.008	0.008	0.02
$q = 5$	—	—	—	$154^{\dagger}$	$132^{\dagger}$	$72^{\dagger}$	88^{*}	66^{*}	54^{*}
$q = 16$	280^{**}	168^{**}	54^{**}	$21^{\dagger}$	$18^{\dagger}$	$12^{\dagger}$	8^{*}	6^{*}	6^{*}
$q = 32$	120^{**}	72^{**}	36^{**}	$7^{\dagger}$	$6^{\dagger}$	$3^{\dagger}$	4^{*}	3^{*}	3^{*}
$q = 64$	80^{**}	48^{**}	18^{**}	$4^{\dagger}$	$3^{\dagger}$	$2^{\dagger}$	2^{*}	2^{*}	1^{*}

TABLE I. Number of round-trip times, N_{R}, required to generate an entangled pair between adjacent nodes with the target fidelity, given initial entangled pairs of fidelity F generated with probability 0.2. Each node contains q matter qubits able to be coupled with deterministic local gates with error rate p_{local}. and l' is the maximum allowed probability that an entangled pair is not generated. * indicates that one round of the [[4,1,2]] code is used, ** indicates that two rounds of the [[4,1,2]] code are used, and † indicates that one round of the [[5,1,3]] code is used.

our communication scheme we must consider the combination of multiplexing, purification, and error correction. The threshold error rate sets the target for the error rate of CZ gates between nodes, which in turn sets the target for the fidelity of entangled pairs after multiplexing and purification, after accounting for the fact that entangled pairs will not always be available. If this target fidelity is met, then we can decrease the logical failure rate arbitrarily by increasing the extent of the cluster state, thereby allowing reliable communication over arbitrarily large distances. How best to allocate resources to multiplexing and purification depends on physical parameters such as the distance between nodes, rate of photon loss, and the accuracy of the local gates. Here, we assume that our nodes are separated by $L =$10 km (corresponding to a round-trip time of flight of $T_R = 0.1$ ms) and that the probability of successfully establishing a raw, entangled pair between two adjacent nodes (before multiplexing) is 20% [54]. This efficiency is beyond present experimental capabilities, but will be useful to illustrate our scheme.

Table I outlines several strategies to generate entangled pairs that meet or exceed the target fidelity using combinations of the [[4,1,2]] and [[5,1,3]] codes in error-detection mode for various values of the initial fidelity, F, and the local gate error rate, p_{local}. As an example, assuming that $F = 0.907$ and $p_{\text{local}} = 0.0008$, with $q = 16$ ($q = 64$) qubits per node, $N_R = 18$ ($N_R = 3$) multiples of the round-trip time are required to generate entanglement of sufficient fidelity using one round of the [[5,1,3]] code. The rate per second at which pairs are generated between adjacent nodes is $R = 1/(N_R T_R) = 10^4/N_R$. For $q = 16$ ($q = 64$), $R \sim 0.5$ kHz ($R \sim 3.3$ kHz). Without additional qubits, schemes that require more than one successful round of purification achieve a significantly lower rate. There is a clear tradeoff between R and q, and more accurate local gates mean that fewer successful rounds of purification are required. Our calculations assume worst-case behaviour, and it is likely that R could be optimized with more careful scheduling. Ultimately, the communication rate is limited by the time to prepare the cluster state, which is $\sim 4T_R$. Quantum memories must be stable over this time, but this requirement is independent of the total communication distance.

Discussion— As our scheme is based on fault-tolerant error correction, it is not surprising that the requirements on the local gates are quite stringent. Superconducting circuits and quantum memories have not achieved this level of accuracy yet, but progress towards this goal has been made. Error rates of quantum gates in superconducting circuits are approaching values as low as one percent [43]. Similarly, quantum memories are being engineered with increasing stability [44, 45]. It may be possible to ease the requirements on these components by increasing the number of qubits per node. Similarly, a scheme based on postselected error correction might lead to a higher threshold [46]. On the other hand, it would be interesting to study a scheme with only a few qubits per node. This would introduce more non-deterministic elements, thereby lowering the threshold [47–49]. Studying these and other schemes will be helpful in the search to strike the correct balance between the complexity and accuracy of the hardware to enable quantum communication over large distances. Lastly, our scheme enables universal quantum computation. In this context, node separation may be shorter and photon loss less severe. This makes our scheme a promising starting point for the study of fault-tolerant quantum computation in hybrid systems based on a resonator-ensemble interaction.

We recently became aware of a closely related manuscript that also proposes the distribution of entanglement using a topological cluster state. See *Long range failure-tolerant entanglement distribution*, Ying Li, Sean Barrett, Thomas Stace, and Simon Benjamin.

Acknowledgements— We acknowledge support from JSPS, MEXT, FIRST, and NICT Japan.

* astephens@unimelb.edu.au

[1] N. Gisin and R. Thew, Nature Photon. **1**, 165 (2007).
[2] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, Phys. Rev. Lett. **81**, 5932 (1998).
[3] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, Phys. Rev. A **59**, 169 (1999).
[4] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, Rev. Mod. Phys. **83**, 33 (2011).
[5] C. H. Bennett *et al.*, Phys. Rev. Lett. **76**, 722 (1996).

[6] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Phys. Rev. A **54**, 3824 (1996).
[7] D. Deutsch *et al.*, Phys. Rev. Lett. **77**, 2818 (1996).
[8] C. H. Bennett *et al.*, Phys. Rev. Lett. **70**, 1895 (1993).
[9] J.-W. Pan, S. Simon, C. Brukner, and A. Zeilinger, Nature **410**, 1067 (2001).
[10] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, Nature **414**, 413 (2001).
[11] B. Zhao, Z.-B. Chen, Y.-A. Chen, J. Schmiedmayer, and J.-W. Pan. Phys. Rev. Lett. **98**, 240502 (2007).
[12] Z.-S. Yuan *et al.*, Nature **454**, 1098 (2008).
[13] L. Hartmann, B. Kraus, H.-J. Briegel, and W. Dür, Phys. Rev. A **75**, 032310 (2007).
[14] S. Perseguers *et al.*, Phys. Rev. A **78**, 062324 (2008).
[15] S. Perseguers, Phys. Rev. A **81**, 012310 (2010).
[16] L. Jiang *et al.*, Phys. Rev. A **79**, 032325 (2009).
[17] A. G. Fowler *et al.*, Phys. Rev. Lett. **104**, 180503 (2010).
[18] A. Grudka *et al.*, arXiv:1202.1016 (2012).
[19] D. Aharonov and M. Ben-Or, Proc. ACM Symp. Th. Comput. **29**, 176 (1998).
[20] W. J. Munro, K. A. Harrison, A. M. Stephens, S. J. Devitt, and K. Nemoto, Nature Photon. **4**, 792 (2010).
[21] R. Raussendorf, J. Harrington, and K. Goyal, Ann. Phys. **321**, 2242 (2006).
[22] R. Raussendorf, J. Harrington, and K. Goyal, New J. Phys. **9**, 199 (2007).
[23] S. D. Barrett and T. M. Stace, Phys. Rev. Lett. **105**, 200502 (2010).
[24] L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, Phys. Rev. Lett. **96**, 070504 (2006).
[25] J. Clarke and F. K. Wilhelm, Nature **453**, 1031 (2008).
[26] F. G. Paauw, A. Fedorov, C. J. P. M. Harmans, and J. E. Mooij, Phys. Rev. Lett. **102**, 090501(2009).
[27] X. Zhu, A. Kemp, S. Saito, and K. Semba, Appl. Phys. Lett. **97**, 102503 (2010).
[28] A. Fedorov *et al.*, Phys. Rev. Lett. **105**, 060503 (2010).
[29] I. Siddiqi *et al.*, Phys. Rev. Lett. **93**, 207002 (2004).
[30] X. Zhu *et al.*, Nature **478**, 221 (2011).
[31] P. Bushev *et al.*, Phys. Rev. B **84**, 060501 (2011).
[32] M. U. Staudt *et al.*, arXiv:1201.1718 (2012).
[33] B. Lauritzen *et al.*, Phys. Rev. Lett. **104**, 080502 (2010).
[34] O. A. Collins, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy, Phys. Rev. Lett. **98**, 060502 (2007).
[35] C. Simon *et al.*, Phys. Rev. Lett. **98**, 190503 (2007).
[36] W. Tittel *et al.*, Laser Photon. Rev. **4**, 244 (2009).
[37] N. Sangouard, R. Dubessy, and C. Simon, Phys. Rev. A **79**, 042340 (2009).
[38] S. N. Dorenbos *et al.*, Appl. Phys. Lett. **93**, 131101 (2008).
[39] R. Van Meter, T. D. Ladd, W. J. Munro, and K. Nemoto, IEEE/ACM Trans. Networking **17**, 1002 (2009).
[40] H. Aschauer, PhD thesis. Ludwig Maximilians Universität, München (2004).
[41] D. Gottesman and I. L. Chuang, Nature **402**, 390 (1999).
[42] C. D. Lorenz and R. M. Ziff, Phys. Rev. E **57**, 230 (1998).
[43] A. Wallraff *et al.*, Phys. Rev. Lett. **95**, 060501 (2005).
[44] M. Steger *et al.*, Science **336**, 1280 (2012).
[45] P. C. Maurer *et al.*, Science **336**, 1283 (2012).
[46] E. Knill, Nature **434**, 39 (2005).
[47] Y. Li, S. D. Barrett, T. M. Stace, and S. C. Benjamin, Phys. Rev. Lett. **105**, 250502 (2010).
[48] K. Fujii and Y. Tokunaga, Phys. Rev. Lett. **105**, 250503 (2010).
[49] Y. Li and S. C. Benjamin, arXiv:1204.0443 (2012).
[50] V. Kolmogorov, Math. Program. Comput. **1**, 43 (2009).
[51] M. Saito and M. Matsumoto, Monte Carlo and Quasi-Monte Carlo Methods **2**, 607 (2006).
[52] We transmit qubits from both sides, leaving a small number of qubits on each side as receivers.
[53] We use a simulator written for this purpose and verified by reproducing the results of Ref. [23]. Our numerical results rely on the Blossom V matching algorithm [50] and the Mersenne Twister pseudorandom number generator [51]. Each data point in Fig. 3 is an average of at least 10^5 trials.
[54] We assume that the rate of loss is the same for connections along all three axes of the cluster state. In practice, spacing perpendicular to the direction of communication may be shorter, meaning that loss due to attenuation will be significantly lower.

Supplementary material

Here we describe two-qubit gates between qubits in the same node. The relevant system comprises two rare-earth qubits (either single atoms or ensembles) prepared in the states $|a\rangle = a_0|0\rangle + a_1|1\rangle$ and $|b\rangle = b_0|0\rangle + b_1|1\rangle$, two gap-tunable flux qubits, $|f_1\rangle$ and $|f_2\rangle$, initially prepared in the ground state, and a resonator, $|r\rangle$, in the vacuum state. The flux qubits are off resonance with both the rare-earth qubits and the resonator. Our initial state is

$$|a\rangle|f_1\rangle|r\rangle|f_2\rangle|b\rangle = |a\rangle|0\rangle|0\rangle|0\rangle|b\rangle. \quad (1)$$

The gate begins by exciting the first flux qubit, $|f_1\rangle$, from the ground state, $|0\rangle$, to the excited state, $|1\rangle$, while it is far off resonance with both its rare-earth qubit and resonator. Then, the flux qubit is quickly moved into resonance with the resonator for a short time to create the entangled state

$$|f_1\rangle|r\rangle = |1\rangle|0\rangle + |0\rangle|1\rangle, \quad (2)$$

then moved out of resonance again. The second flux qubit is brought into resonance with the resonator until the SWAP gate exchanges the excitation between the resonator and flux qubit (this takes twice the entangling time), then moved out of resonance again. After a single-qubit rotation on the first qubit (which can be done in parallel with the previous step), we have generated the entangled state of the flux qubits

$$|f_1\rangle|f_2\rangle = |0\rangle|0\rangle + |1\rangle|1\rangle, \quad (3)$$

with no population in the resonator.

Next, the state of the flux qubits is used to perform an effective CZ gate between the rare-earth qubits. First, both flux qubits are brought into resonance with their respective rare-earth qubits to allow a dispersive interaction. This interaction generates a CZ between each pair of flux qubit and rare-earth qubit, so a four-qubit entangled state is prepared. Then, a Hadamard rotation on each flux qubit and measurements in the computational basis project the rare-earth qubits to the state

$$a_0|0\rangle b_0|0\rangle + a_0|0\rangle b_1|1\rangle + a_1|1\rangle b_0|0\rangle - a_1|1\rangle b_1|1\rangle, \quad (4)$$

up to a known correction. That is, an effective CZ gate has been performed between the two rare-earth qubits. This operation can be completed in approximately 150 ns, assuming strong enough coupling between the rare-earth qubits and the flux qubits.